Estimating the probability distribution of the future exchange rate from opti...
Allan M Malz
Journal of Derivatives; Winter 1997; 5, 2; ABI/INFORM Global
pg. 18

ESTIMATING THE PROBABILITY DISTRIBUTION OF THE FUTURE EXCHANGE RATE FROM OPTION PRICES

Allan M. Malz
is with the Markets Group of the Federal Reserve Bank of New York.

This article describes a method of extracting the risk-neutral probability distribution of future exchange rates from option prices. In foreign exchange markets, interbank option pricing conventions facilitate reliable inferences about risk-neutral probability distributions with a small amount of readily available information.

The risk-neutral probability distribution of the future exchange rate provides investors and market analysts with an important tool for gauging market sentiment.

Market participants need to understand market sentiment and keep updated on its changes. Central banks responsible for maintaining the stability of the exchange markets are more concerned about an exchange rate move when it is associated with great uncertainty about future moves or anxiety about their direction. Investors make better-informed judgments when they compare their own view of future exchange rate developments to the market consensus.

Market beliefs about a future asset price and market attitudes toward the risks of particular asset price realizations can be compactly summarized in the risk-neutral probability distribution, the distribution implied by the observed market prices of the asset and its derivatives. The risk-neutral distribution may differ from the market's subjective probability distribution if market prices include premiums for risk and liquidity.

For example, the forward foreign exchange rate can be thought of as the mathematical expectation of the future spot exchange rate, with the probability weights drawn from the risk-neutral distribution. The forward rate may or may not correspond to the true market consensus forecast of the future spot rate, but it does represent the price at which market participants can tailor their foreign exchange exposures, laying off or taking on foreign exchange risk.

Similarly, the market price of a European currency call option can be thought of as the risk-neutral mathematical expectation of its payoff at maturity, discounted to the present. In other words, it is the expected value of the future exchange rate less the exercise price, with the probability weights drawn from the risk-neutral distribution. If the market is eager to insure against the risk of currency appreciation, calls with high exercise prices might be more expensive than if they represented a fair bet on appreciation. In that case, the risk-neutral probability of an appreciation might be higher than the market's subjective probability.

Option prices provide much embedded information on the risk-neutral distribution. In fact, if prices of calls and puts with the same expiration but many different exercise prices could be simultaneously observed, one could trace out the entire risk-neutral distribution, since the discounted risk-neutral density function of the asset price equals the second derivative of the call option price with respect to the exercise price:[1]

$$\frac{\partial^2 c(t, X, T)}{\partial X^2} = e^{-r\tau}\pi(X) \qquad (1)$$

where c(t, X, T) represents the observed time t market price of a European call option on an asset with an exercise price of X expiring at time T.

This suggests that one could extract the probability distribution of the asset price by taking second differences of simultaneously observed call option prices with different exercise prices. Unfortunately, this is difficult to apply, because it requires, in principle, options with a continuous — or at least closely spaced — series of exercise prices. In practice, not many option contracts on a given asset with different exercise prices and a given maturity trade in the marketplace at any one time, so implementing this approach requires some sort of interpolation between observed option prices.[2]

An alternative solution to the "shortage" of option data adds structure by assuming that the risk-neutral probability distribution of the future asset price belongs to a particular parametric family other than the lognormal, for example a jump diffusion. The distributional assumption leads to option pricing formulas in which the parameters of the distribution appear. The parameters can be estimated by fitting the option data to the pricing formulas.[3]

To date, such techniques for extracting risk-neutral distributions from option prices have been anything but straightforward; they require much data and are computationally difficult. This article describes a simplified interpolation method for extracting the risk-neutral distribution of the future exchange rate from option prices. The simplification is made possible by relying on price data from the over-the-counter option markets rather than centralized exchanges. Dealer conventions in over-the-counter markets permit greater accuracy and ease of computation with a small amount of data that are readily available from public sources and are accessible to investors, nonfinancial corporations, and central banks.

As is well known, forward asset prices are poor indicators of the future behavior of cash prices, since they embody little more information about future asset prices than the asset's cash price. Observation of the entire risk-neutral probability distribution of the future exchange rate, rather than just the first moment, gives investors and market analysts an important additional tool with which to gauge market sentiment.

I. THE OVER-THE-COUNTER CURRENCY OPTION MARKETS

The language and conventions that traders in the over-the-counter currency option markets use are borrowed from the Black-Scholes model, even though traders are fully aware that the model is at best an approximation.[4] The model, which assumes that the spot exchange rate follows geometric Brownian motion, results in the familiar formula for the price of a European call:

$$v(S_t, \tau, X, \sigma, r, r^*) =$$

$$e^{-r^*\tau} S_t \Phi\left[\frac{\ln\left(\frac{S_t}{X}\right) + \left(r - r^* + \frac{\sigma^2}{2}\right)\tau}{\sigma\sqrt{\tau}}\right] -$$

$$e^{-r\tau} X \Phi\left[\frac{\ln\left(\frac{S_t}{X}\right) + \left(r - r^* - \frac{\sigma^2}{2}\right)\tau}{\sigma\sqrt{\tau}}\right] \qquad (2)$$

where S_t represents the exchange rate, $\tau \equiv T - t$ the time to expiration, X the exercise price, σ a parameter governing the variance of the exchange rate, r and r^* the domestic and foreign interest rates, and Φ represents the standard cumulative normal distribution function. The probability distribution implied by the Black-Scholes formulas is the lognormal distribution, as can be seen by calculating the second derivative of Equation (2) with respect to X.

Given an observed call price c(t, X, T) and the observed or contractually specified values of S_t, τ, X, r_t, and r_t^*, the equation

$$c(t, X, T) = v(S_t, \tau, X, \sigma_t, r_t, r_t^*) \qquad (3)$$

can be solved for the unique implied volatility σ_t corresponding to c(t, X, T).

Dealers in over-the-counter markets quote option prices in terms of implied volatility, in units called *vols*. When a deal is struck, they translate the agreed price from vols to currency units by plugging current market data and the agreed exercise price, maturity, and option price in vols into the right-hand side of Equation (3). Exhibit 1A shows how these calculations are carried out.

Much over-the-counter option trading is in at-the-money forward options, for which the exercise price is set equal to the current forward exchange rate. When in- or out-of-the-money options are dealt, exercise prices are not expressed in currency units, but in terms of delta, the rate of change of the Black-Scholes call option value with respect to the spot exchange rate:

$$\delta_v(S_t, \tau, X, \sigma, r, r^*) \equiv \frac{\partial v(S_t, \tau, X, \sigma, r, r^*)}{\partial S_t} =$$

$$e^{-r^*\tau}\Phi\left[\frac{\ln\left(\frac{S_t}{X}\right) + \left(r - r^* + \frac{\sigma^2}{2}\right)\tau}{\sigma\sqrt{\tau}}\right] \qquad (4)$$

The call delta always lies on $[0, e^{-r^*\tau}]$. The delta of a put is $\delta_v(S_t, \tau, X, \sigma, r, r^*) - e^{-r^*\tau}$. Often, exercise prices are set so that delta is equal to a round number like 25%. If a customer buys, say, a 25-delta dollar-mark call, the exercise price is calculated by setting the left-hand side of Equation (4) equal to 0.25 and solving for X. The call delta declines monotonically as exercise price rises — and the option goes farther out of the money — so dealers can readily find the unique exercise price corresponding to a given delta, and vice versa. Exhibit 1B illustrates these calculations.

Dealers usually express the delta of an option as a percent, not a decimal, and omit the minus sign

EXHIBIT 1
OPTION PRICE QUOTE CONVENTIONS

PANEL A. IMPLIED VOLATILITY AND OPTION PRICE IN CURRENCY UNITS

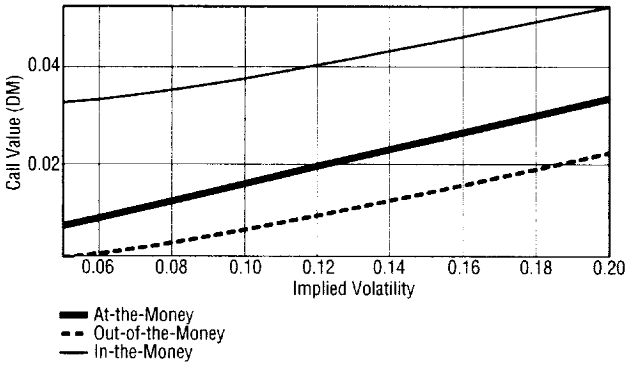

PANEL B. CALL OPTION DELTA AND EXERCISE PRICE

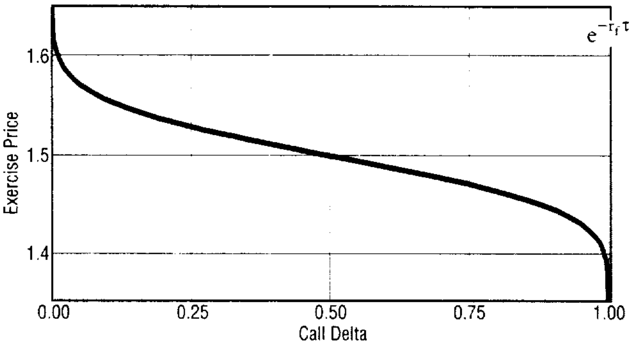

Dollar call against the German mark, spot exchange rate DM 1.50, maturity one month, implied volatility 10%. Interest rates: domestic 3%, foreign 5%.

in discussing put option deltas, so a "25-delta put" refers to a put with a delta of –0.25. The delta of an at-the-money forward call (put) option is close, but not exactly equal to 0.5 (–0.5).

II. THE VOLATILITY SMILE

The Black-Scholes model assumes that all options on the same currency have identical implied volatilities, regardless of time to maturity and moneyness. However:

- Options with the same exercise price but different maturities often have different implied volatilities, giving rise to a term structure of implied volatility. A rising term structure indicates that market participants expect short-term implied volatility to rise or that they are willing to pay more for protection against near-term exchange volatility.
- Out-of-the-money options on currencies with flexible exchange rates often have higher implied volatilities than at-the-money options, indicating that the market perceives exchange rates to be leptokurtotic; that is the risk-neutral likelihood of large exchange rate moves is greater than is consistent with the lognormal distribution.
- Out-of-the-money call options often have implied volatilities that differ from those of equally out-of-the-money puts, indicating that the market perceives the distribution of future exchange rates to display skewness or that market participants are willing to pay more for protection against sharp currency moves in one direction than in the other.

The latter two phenomena constitute the "volatility smile," so-called because of the characteristic shape of the plot of implied volatilities of options of a given maturity against delta. The volatility smile is evidence not only that the Black-Scholes model does not hold exactly, but also that dealers are perfectly aware of this. The market uses the Black-Scholes terminology nonetheless, raising or lowering the prices, and therefore implied volatilities, of options with different exercise prices to equate their supply and demand.

We can summarize the observed volatility smile at time t for call options with maturity τ in the function $\sigma_X(t, X, T)$, which gives the implied volatility corresponding to each exercise price for which an option price can be observed. It is related to observed option prices and to the Black-Scholes option pricing formulas by

$$c(t, X, T) = v[S_t, \tau, X, \sigma_X(t, X, T), r_t, r_t^*]$$

so implied volatilities are calculated the same way when there is a volatility smile as when implied volatility is constant. Because options with only a handful of different exercise prices trade actively at any moment, $\sigma_X(t, X, T)$ is not a continuous function.

As noted, over-the-counter currency option dealers often express exercise prices in terms of delta rather than currency units. The market-implied volatility schedule $\sigma_X(t, X, T)$ can be expressed in terms of delta as $\sigma_\delta(t, \delta, T)$. The two schedules $\sigma_\delta(t, \delta, T)$ and $\sigma_X(t, X, T)$ embody the same information about observed implied volatilities, and are related to one another in a straightforward if non-linear fashion. The exercise price, for example, of a 25-delta call option, which we denote $X_t^{(0.25)}$, can be calculated by setting

$$\delta_v[S_t, \tau, X_t^{(0.25)}, \sigma_\delta(t, 0.25, T), T, r_t, r_t^*] = 0.25$$

Only a change of units is involved, so $\sigma_\delta(t, 0.25, T) \equiv \sigma_X(t, X_t^{(0.25)}, T)$. This identity will hold, of course, for any delta:

$$\sigma_\delta(t, x, T) \equiv \sigma_X(t, X_t^{(x)}, T), \quad x \in [0, e^{-r^*\tau}] \tag{5}$$

where $X_t^{(x)}$ satisfies

$$\delta_v(S_t, \tau, X_t^{(x)}, \sigma_\delta(t, x, T), T, r_t, r_t^*) = x, \quad x \in [0, e^{-r^*\tau}] \tag{6}$$

Note also that the implied volatilities of a call with a delta of x and a put with a delta of $x - e^{-r^*\tau}$ are identical, so we can readily include the observed implied volatilities of puts in the schedule $\sigma_\delta(t, \delta, T)$.

Over-the-counter dealers usually express option prices in terms of implied volatility, so $\sigma_\delta(t, \delta, T)$ summarizes the market's price schedule at time t of options on a given currency pair and with a given maturity, but different exercise prices. We can translate the quote on a call option with a delta of x from implied volatility units into currency units via

$$c(t, X_t^{(x)}, T) \equiv$$

$$v(S_t, \tau, X_t^{(x)}, \sigma_\delta(t, 0.25, T), r_t, r_t^*) \qquad (7)$$

where $X_t^{(x)}$ is calculated from Equation (6).

Dealers often sell options in combinations. The *straddle* is a combination of a put and a call with identical exercise prices, usually at the money forward. Two combinations, the *strangle* and the *risk reversal*, involve two out-of-the-money options with the same delta and contain most of the information about the shape of the volatility smile. The strangle is a position consisting of a long out-of-the-money put and call. The risk reversal consists of a long call and a short put.

Exhibit 2 illustrates the payoffs at maturity on these instruments. Strangles and risk reversals generally involve 25-delta calls and puts, but 15-delta strangles and risk reversals are becoming more common in interdealer trading. Quotes for strangles and risk reversals with other deltas are easily obtained over the counter.

The prices of these option combinations are expressed in vols rather than currency units in over-the-counter market parlance. The implied volatility of the put component of an at-the-money forward straddle is identical to that of the call and is referred to (somewhat loosely) as the straddle or at-the-money forward volatility. It indicates the general level of implied volatility for options of a given maturity on a particular currency pair.

In a 25-delta strangle, the dealer sells or buys both out-of-the-money options from the counterparty. Dealers usually quote strangle prices by stating the implied volatility — the strangle volatility — at which they buy or sell both options. For example, the dealer might quote the selling price as 14.6 vols, meaning a 25-delta call and a 25-delta put sold at an implied volatility of 14.6 vols each. Dealers generally record strangle prices as the spread of the strangle volatility over the at-the-money forward volatility. If market participants were convinced that exchange rates move lognormally, the out-of-the-money options would have the same implied volatility as at-the-money options, and strangle spreads would be zero. Strangles, then, indicate the degree of curvature of the volatility smile.

EXHIBIT 2

STRADDLE, STRANGLE, AND RISK REVERSAL PAYOFFS AT MATURITY

PANEL A. STRADDLE PAYOFF

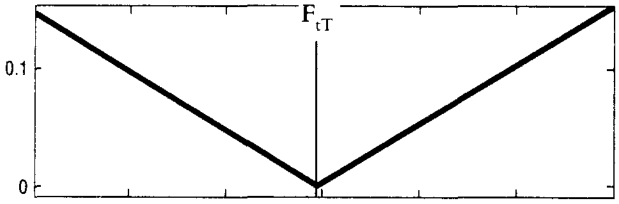

PANEL B. STRANGLE PAYOFF

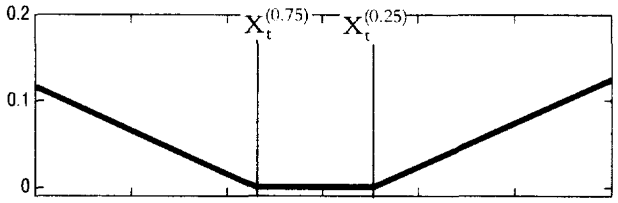

PANEL C. RISK REVERSAL PAYOFF

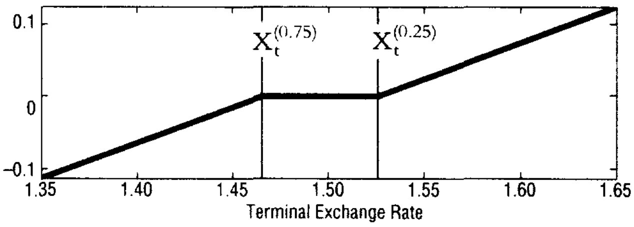

Exercise price DM 1.50; payoff in marks. F_{tT} refers to the one-month forward exchange rate, $X_t^{(0.25)}$ to the exercise price price of a 25-delta option, and $X_t^{(0.75)}$ to the exercise price of a 75-delta option.

In a risk reversal, the dealer exchanges one of the options for the other with the counterparty. Because the put and the call are generally not of equal value, the dealer pays or receives a premium for exchanging the options, quoted as the implied volatility differential for exchanging a 25-delta call for a 25-delta put. For example, if dollar-mark is strongly expected to fall (dollar depreciation), an options dealer might quote dollar-mark risk reversals as: "one-month 25-delta risk reversals are 0.8 at 1.2 mark calls over." This means the dealer stands ready to pay a net premium of 0.8 vols to buy a 25-delta mark call and sell a 25-delta mark put against the dollar, and charges a net premium of 1.2 vols to sell a 25-delta mark call and buy a 25-delta mark put.

Denoting the 25-delta and 75-delta call volatilities by $\sigma_t^{(0.25)} \equiv \sigma_\delta(t, 0.25, T)$ and $\sigma_t^{(0.75)} \equiv \sigma_\delta(t, 0.75, T)$ and the strangle price, risk reversal price, and at-the-money volatility by str_t, rr_t, and atm_t, the midpoint of the time t strangle price can be expressed as $str_t = 0.5(\sigma_t^{(0.75)} + \sigma_t^{(0.25)}) - atm_t$. The risk reversal price is $rr_t = \sigma_t^{(0.25)} - \sigma_t^{(0.75)}$. Using these definitions, the market quotes for the strangle price, risk reversal price, and at-the-money volatility can be solved for $\sigma_t^{(0.25)}$ and $\sigma_t^{(0.75)}$:[5]

$$\sigma_t^{(0.25)} = atm_t + str_t + 0.5rr_t$$

$$\sigma_t^{(0.75)} = atm_t + str_t - 0.5rr_t \qquad (8)$$

To close a straddle deal, the counterparties use the Black-Scholes formulas, Equation (2) and its analogue for puts, to translate the straddle price in vols into currency units. To close a strangle or risk reversal deal, the exercise prices of the individual components must first be set, which in turn requires the counterparties to agree on $\sigma_t^{(0.25)}$ and $\sigma_t^{(0.75)}$, the implied volatilities of the 25-delta call and the 25-delta put. These are substituted into Equation (4), which is set equal to 0.25 or 0.75 to solve for the exercise prices. The counterparties then translate the price in vols into currency units using Equation (7).[6]

III. ESTIMATION METHOD

Recall that, in the spirit of Equation (1), we want to estimate the risk-neutral probability distribution by second-differencing the current prices of call options with respect to their exercise prices. Since only a handful of options with different exercise prices trade simultaneously, we need to approximate c(t, X, T) by interpolating a smooth function through the available prices.

It is difficult to interpolate using exchange-traded currency option data. Liquidity and transaction volume in the currency options market are centered in the over-the-counter market.[7] The exchanges' settlement price data almost invariably include prices of barely traded or untraded options that do not accurately reflect the state of the market for those exercise prices, and frequently even appear to violate no-arbitrage conditions. There is no absolute way to identify such unrepresentative prices. The few options in which there is a significant open interest are usually bunched together near the current spot rate, so only a small portion of the exercise price axis is covered, making inferences about moments difficult and prone to error. What seems like an abundance of simultaneous indicative prices for options with different exercise prices on the exchanges is thus deceptive.

We therefore use over-the-counter currency option data. Simultaneous public price data (via Reuters) for one-month options with the exercise prices $X_t^{(0.25)}$, $X_t^{(0.50)}$, and $X_t^{(0.75)}$ are currently available.[8] These three prices are well-spaced over the exercise price axis, permitting an accurate interpolation to be made.

We do not approximate c(t, X, T) directly, but rather the market volatility smile schedule $\sigma_\delta(t, \delta, T)$, since the over-the-counter markets quote currency option prices in vols and set exercise prices in terms of delta. The next step is to get from delta–volatility space to exercise price–volatility space by translating the approximated version of $\sigma_\delta(t, \delta, T)$, which relates implied volatility to delta, into $\sigma_X(t, X, T)$, which relates implied volatility to exercise price, through Equation (8). Finally, we use Equation (7) to get from exercise price–volatility space to an approximated c(t, X, T), in exercise price–currency unit space.

The first step is to approximate $\sigma_\delta(t, \delta, T)$ by the function

$$\hat{\sigma}_\delta(\delta) = b_0 atm_t + b_1 rr_t(\delta - 0.50) + b_2 str_t(\delta - 0.50)^2$$

Mathematically, this functional form can be interpreted as a second-order Taylor approximation to the volatility smile or a spline approximation to the three data points with parabolic endpoints. It is easy to solve for the parameters b_i, $i = 0, 1, 2$, by imposing the condition that the at-the-money volatility and the risk reversal and strangle prices lie exactly on $\hat{\sigma}_\delta(\delta)$ (see Appendix B). This forces our approximating function $\hat{\sigma}_\delta(\delta)$ to exactly equal the observed market implied volatilities for $\delta = (0.25, 0.50, 0.75)$ and helps guarantee that it is a good approximation. The resulting functional form is:

$$\hat{\sigma}_\delta(\delta) = atm_t - 2rr_t(\delta - 0.50) + 16str_t(\delta - 0.50)^2 \qquad (9)$$

The minimum of this function is found at $\delta = 0.50 + rr_t/16str_t$, which is generally quite close to 0.50.

Having the observed option prices lie on the interpolated volatility smile avoids imposing a particular hypothesis about the exchange rate's probability distribution, while minimizing the errors in the interpolation. The functional form (9) is the simplest one that captures the information about the smile that the three option prices express. The at-the-money volatility gives the general level of implied volatility; it is a "measure of location" for the volatility smile. The risk reversal price indicates the skewness in the smile, and the strangle price indicates the degree of curvature of the smile.

Exhibit 3A illustrates the interpolation with a fairly typical example for a floating-rate currency under selling pressure: $atm_t = 0.10$, $rr_t = -0.015$, and $str_t = 0.005$.

Once we have approximated the market's schedule of implied volatilities by $\hat{\sigma}_\delta(\delta)$, we must find the implied volatility corresponding to each exercise price rather than each delta. Delta itself is a function of the implied volatility, so we substitute Equation (4) into (9), and solve for σ as a function of X:

EXHIBIT 3
VOLATILITY SMILE

PANEL A. IMPLIED VOLATILITY AS A FUNCTION OF DELTA

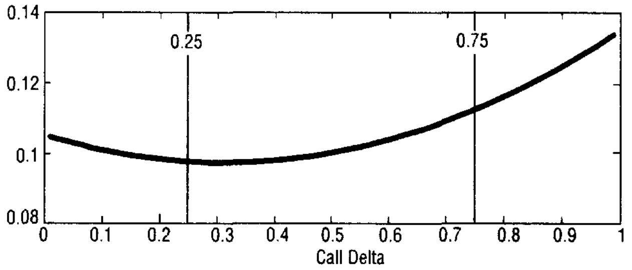

PANEL B. IMPLIED VOLATILITY AS A FUNCTION OF EXERCISE PRICE

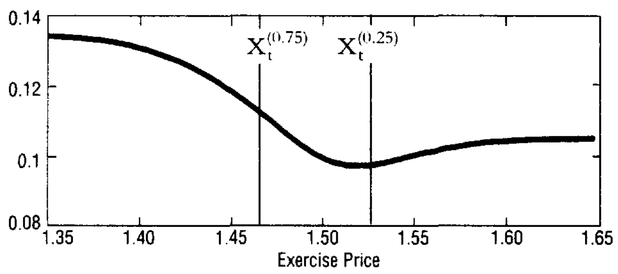

PANEL C. RISK-NEUTRAL DENSITY

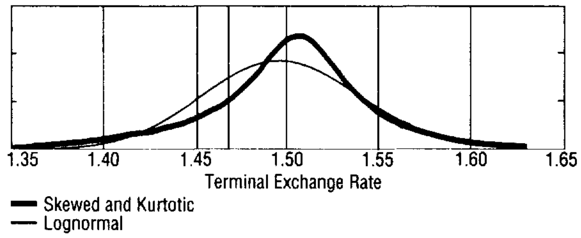

$$0 = \sigma - atm_t + 2rr_t \times \left[e^{-r^*\tau}\Phi\left[\frac{\ln\left(\frac{S_t}{X}\right) + \left(r - r^* + \frac{\sigma^2}{2}\right)\tau}{\sigma\sqrt{\tau}}\right] - 0.50\right] -$$

$$16\text{str}_t\left[e^{-r^*\tau}\Phi\left[\frac{\ln\left(\frac{S_t}{X}\right)+\left(r-r^*+\frac{\sigma^2}{2}\right)\tau}{\sigma\sqrt{\tau}}\right]-0.50\right]^2 \tag{10}$$

There is no closed-form solution to Equation (10), which we refer to as the *implied volatility function* and denote $\sigma_X(X)$. It is easy to solve numerically.[9]

Exhibit 3B displays the volatility function for the example. Its appearance may be surprising, since it is concave to the origin for deep in- and out-of-the-money options.[10] We show below that this property of volatility smiles is consistent with the no-arbitrage restrictions on option prices.

Substituting the volatility function, Equation (10), into the Black-Scholes call value function, Equation (2), yields a generalized Black-Scholes formula in which the implied volatility depends on the exercise price and can thus be substituted out:

$$\begin{aligned}\hat{v}(X) &\equiv v(S_t, \tau, X, \sigma_X(X), r, r^*)\\ &= e^{-r^*\tau}S_t\Phi\left[\frac{\ln\left(\frac{S_t}{X}\right)+\left(r-r^*+\frac{\sigma_X(X)^2}{2}\right)\tau}{\sigma_X(X)\sqrt{\tau}}\right]-\\ &\quad e^{-r\tau}X\Phi\left[\frac{\ln\left(\frac{S_t}{X}\right)+\left(r-r^*-\frac{\sigma_X(X)^2}{2}\right)\tau}{\sigma_X(X)\sqrt{\tau}}\right]\end{aligned} \tag{11}$$

We can now calculate the estimates of the cumulative distribution function, denoted $\hat{\Pi}(X)$, which estimates the probability that the future exchange rate will be less than or equal to X, and the probability density function, denoted $\hat{\pi}(X)$. Twice differentiating Equation (11) numerically with respect to X and multiplying by $e^{-r\tau}$ yields the estimated probability density function.

The easiest way to carry this out is by difference quotients: 1) For each X, calculate $\sigma_X(X)$ numerically using Equation (10); 2) the estimated cumulative distribution function at that point is

$$\hat{\Pi}(X) = e^{-r\tau}\left[\frac{\hat{v}(X)-\hat{v}(X-\Delta X)}{\Delta X}+1\right]$$

and the estimated probability density function is

$$\hat{\pi}(X) = \frac{\hat{\Pi}(X)-\hat{\Pi}(X-\Delta X)}{\Delta X}$$

where ΔX is the step size; 3) repeat for each X to draw the entire cumulative distribution or density function.[11]

Exhibit 3C illustrates the results for the example, comparing the risk-neutral distribution, which displays substantial kurtosis and negative skewness, with the lognormal distribution implied by the at-the-money volatility under the Black-Scholes assumption that the exchange rate follows a geometric Brownian motion.

This procedure can be used to calculate the higher moments of the risk-neutral distribution. The risk-neutral standard deviation is usually quite close to the straddle price — that is, the at-the-money forward implied volatility — but the risk-neutral coefficients of kurtosis and skewness can differ substantially from the near-zero values they take on in the Black-Scholes model.[12]

The data required to carry out the procedure will come from two sources. The straddle, strangle, and risk reversal prices can be obtained by asking a foreign exchange dealer or broker. Several Reuters screens also carry the data. Several dealers or screens can be polled to get a representative reading of market levels. The remaining market prices, the spot exchange rate and short-term interest rates, are available from any electronic data feed service.

IV. ACCURACY

The accuracy of this method, which we will refer to as the volatility function method, depends on whether the interpolated volatility smile is a good approximation to actual market prices. How close are the interpolated smile represented by Equation (9) and the call prices it generates to the implied volatilities and observed prices of call options with deltas other than 25%, 50%, and 75%? What impact do the interpolation errors have on estimates of the risk-neutral distribution?

We will explore these issues by comparing the values generated by the volatility function technique with actual dealer quotes and with the values generated by an alternative technique, that of Shimko [1993].

Exhibit 4 compares the interpolation given by Equation (9) with a "ladder" of one-month dollar-mark and dollar-yen call option volatilities with $\delta = [0.05, 0.10, ..., 0.95]$ provided by two dealers during the New York afternoon of March 6, 1997. These indicative levels are not generally available to the public.[13]

The interpolation, of course, coincides with the quoted 25-, 50-, and 75-delta call volatilities. Because the 25-, 50-, and 75-delta options anchor the interpolated smile, Equation (9) is very close to actual market implied volatilities for deltas between about 15 and 85. For options with deltas outside that range, so-called wing options, trading is thinner, and dealers often have rather different quotes for, say, 10- and 90-delta options, making it more difficult to establish "the" market level.

Dealers set implied volatilities for less actively traded deltas by combining an interpolation technique with judgment, depending on their sense of the market and their wish to position their option portfolios in particular ways. The two sets of implied volatility quotes represent the range of approaches to pricing wing options; Dealer I prefers here to set low, and Dealer II sets somewhat higher implied volatilities.

The differences in wing option volatilities among dealers are related to differences in views about the likelihood of very large moves in exchange rates, in hedging needs, or in anticipated order flows for these options. Since the price differences in currency units are marginal, and dealers do not typically take on large net positions, dealers quoting high do not buy large quantities of options from dealers quoting low. Rather, their wing option quotes express a general posture.

The small differences in implied volatilities for deltas close to zero or $e^{-r^*\tau}$ imply only small differences in option prices expressed in currency units, since for very high and low exercise prices, the option vega — the rate of change of the option's value with respect to the volatility parameter — is very low (see Appendix C for details). Exhibits 5 and 6 illustrate by converting the implied volatilities into dollar call price quotes in marks or yen. The largest difference between a dealer's quotes and the interpolation, for 5-delta dollar calls, is 3/100 of one pfennig, or 27%. The remaining interpolation values are almost identical to the actual quotes.

Small errors in interpolation have in general only a small impact on estimates of the risk-neutral distribution. Exhibit 7 illustrates by comparing estimates of the risk-neutral density generated by the volatility function technique to estimates generated by passing a cubic spline function through the full ladder of dealer indications for one-month dollar-mark and dollar-yen calls on March 6, 1996.

I fit a cubic spline with linear endpoints in call delta-implied volatility space. The spline interpolation generates an estimate of the call value that is anchored to dealer indications for nineteen rather than only three exercise prices. I estimate the risk-neutral distribution and density functions using difference quotient approximations to the first and second derivatives of the call value with respect to the exercise price.[14]

The volatility function and spline estimates are also compared to estimates using the at-the-money implied volatility under the Black-Scholes assumption of lognormality. The volatility function and spline estimates are fairly close to one another, even in the tails of the distributions, and are different from the lognormal estimates.

Another comparison can be made with the implied volatilities and option prices generated by Shimko's [1993] technique, which, like ours, proposes interpolating implied volatilities rather than option prices, but, in contrast to ours, interpolates them as a function of exercise price rather than delta. Shimko fits a quadratic function of exercise price to the implied volatility data by linear least squares:

$$\sigma_X(t, X_i, T) = a_0 + a_1 X_i + a_2 X_i^2 + u_i \quad (12)$$

EXHIBIT 4

DEALER QUOTES AND INTERPOLATED IMPLIED VOLATILITIES

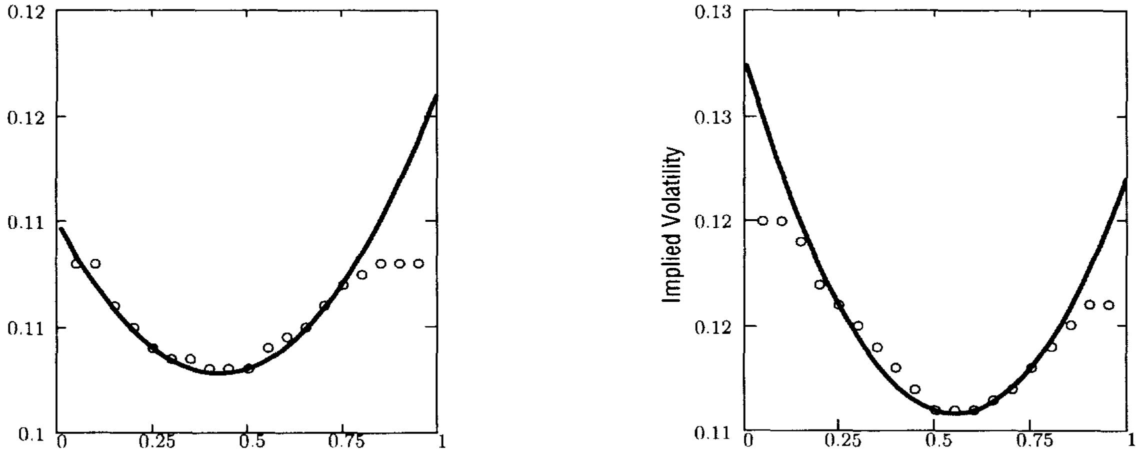

PANEL A. DOLLAR-MARK, DEALER I, MARCH 6, 1997

PANEL C. DOLLAR-YEN, DEALER I, MARCH 6, 1997

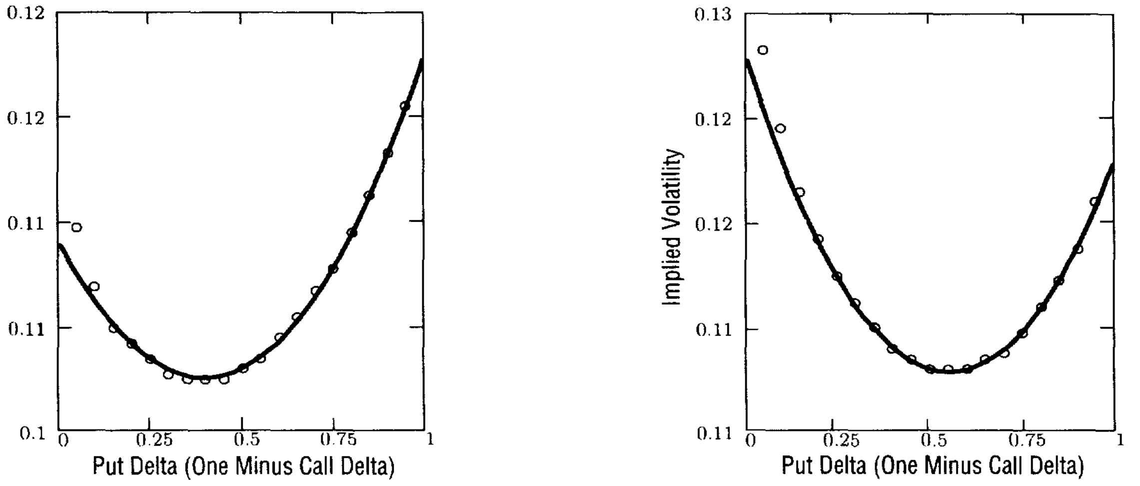

PANEL B. DOLLAR-MARK, DEALER II, MARCH 6, 1997

PANEL D. DOLLAR-YEN, DEALER II, MARCH 6, 1997

Circles represent dealer quotes, and solid lines represent values derived by interpolation.

The $(\sigma_X(t, X_i, T), X_i)_{i=1}^n$ are pairs of implied volatilities and the corresponding exercise prices, and the $(u_i)_{i=1}^n$ are error terms.

Since there are three parameters to estimate, the fit to our three data points is exact: $u_1 = u_2 = u_3 = 0$.[15] A call price with any given exercise price is calculated by substituting the interpolated implied volatility corresponding to that exercise price into Equation (2).

Exhibit 8, which is based on the example in Exhibits 2 and 3, illustrates the differences between the Shimko and the volatility function techniques. The two interpolations coincide for most deltas. One key difference is that Shimko's technique results in

EXHIBIT 5
DEALER QUOTES AND INTERPOLATED CALL OPTION VOLATILITIES

PANEL A. DOLLAR-MARK, DEALER I, MARCH 6, 1997

PANEL C. DOLLAR-YEN, DEALER I, MARCH 6, 1997

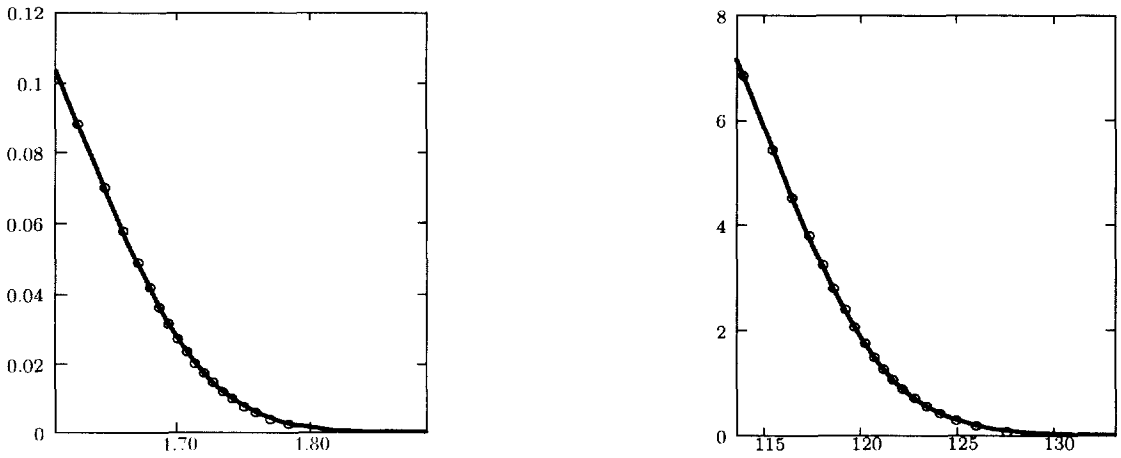

PANEL B. DOLLAR-MARK, DEALER II, MARCH 6, 1997

PANEL D. DOLLAR-YEN, DEALER II, MARCH 6, 1997

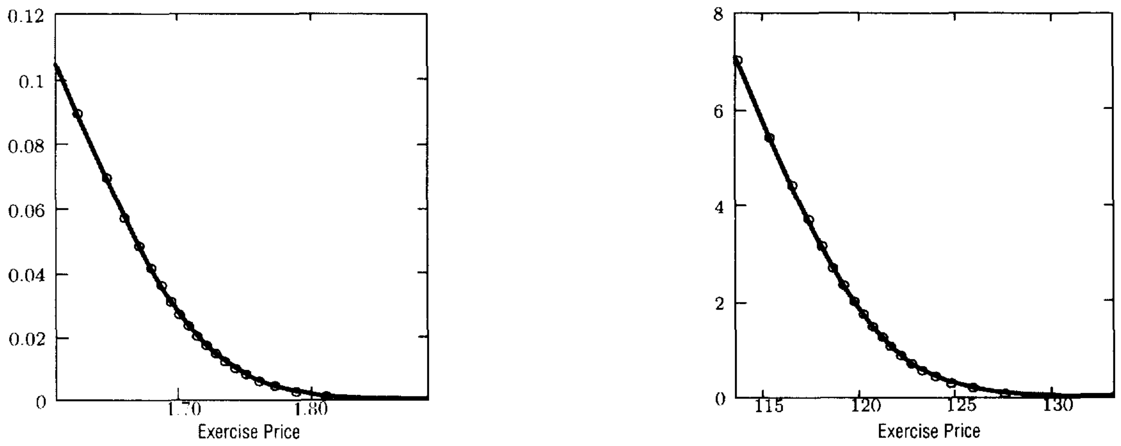

Circles represent dealer quotes, and solid lines represent values derived by interpolation.

significantly higher volatility estimates for wing options and is likely to overestimate them considerably. Evidence for this is provided by Exhibit 6, which displays the call prices generated by Shimko's interpolation function, and shows that they are as much as 82% higher than actual dealer quotes for 5-delta options, a far bigger error than that of the volatility function technique.

The accuracy advantage of the volatility function, with its concave-to-the-origin ranges, should now be evident. Work by Heynen [1994] and Taylor and Xu [1994] indicates that volatility smiles in the currency and stock index option markets can be approximated by quadratic functions of the exercise price or of the ratio of the exercise price to the forward price. These studies are based on exchange-traded option data, however, for which there are relatively few and infrequent observations on wing options. A

EXHIBIT 6

ACCURACY OF ALTERNATIVE INTERPOLATION METHODS

	DEALER I			DEALER II		
DELTA	QUOTE	VOLATILITY FUNCTION	SHIMKO	QUOTE	VOLATILITY FUNCTION	SHIMKO
DOLLAR-MARK CALL OPTION PRICES (IN PFENNIGS)						
0.05	0.11	0.14	0.20	0.12	0.12	0.19
0.10	0.25	0.28	0.32	0.26	0.26	0.31
0.25	0.78	0.78	0.78	0.78	0.78	0.78
0.40	1.45	1.44	1.44	1.45	1.45	1.45
0.60	2.70	2.70	2.70	2.69	2.69	2.69
0.75	4.16	4.16	4.16	4.15	4.15	4.15
0.90	6.97	6.96	6.98	6.91	6.90	6.92
0.95	8.78	8.78	8.81	8.91	8.90	8.94
DOLLAR-MARK CALL OPTION PRICES (IN YEN)						
0.05	0.08	0.10	0.14	0.08	0.08	0.12
0.10	0.19	0.20	0.23	0.19	0.19	0.21
0.25	0.58	0.58	0.58	0.56	0.56	0.56
0.40	1.09	1.09	1.09	1.06	1.06	1.06
0.60	2.09	2.08	2.07	2.01	2.02	2.01
0.75	3.27	3.27	3.27	3.17	3.17	3.17
0.90	5.45	5.46	5.50	5.43	5.42	5.45
0.95	6.86	6.87	6.92	7.03	7.03	7.07

Indicative levels, March 6, 1997.

EXHIBIT 7

ESTIMATED RISK-NEUTRAL DENSITY FUNCTIONS: VOLATILITY FUNCTION AND SPLINE INTERPOLATIONS

PANEL A. DOLLAR-MARK, DEALER I, MARCH 6, 1997

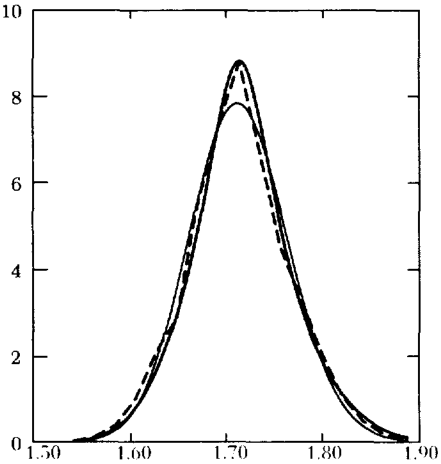

PANEL B. DOLLAR-MARK, DEALER II, MARCH 6, 1997

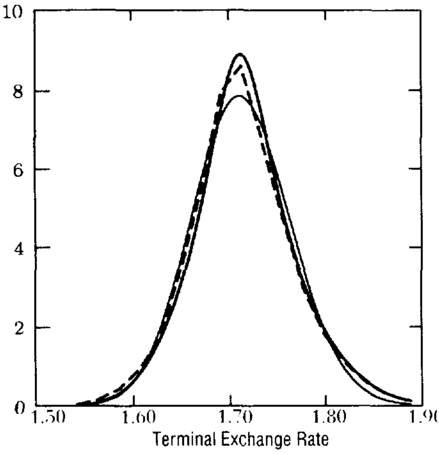

more complex functional form that permits the smile slope to taper off to zero as delta approaches zero or one might well be found to fit the data as well as or better than the simple quadratic. A dealer interpolating in implied volatility–exercise price space might or might not make the interpolation strictly convex even for high and low deltas, but as we have seen, this would translate into small differences in call prices, since even a large change in volatility leads to a small change in price for the wing options.

Exhibit 8 displays estimates of the risk-neutral density using our technique and Shimko's, and a simple estimate assuming lognormality. The two techniques that take the volatility smile into account are roughly similar to one another and quite different from the simple lognormal approach. As might be expected, the Shimko density displays somewhat greater kurtosis than the density generated by the volatility function.

Another key difference is that, for wing options, Shimko's technique violates the no-arbitrage conditions in Appendix C. The second derivative of Shimko's estimated volatility smile is $a_1 + 2a_2X$, which does not disappear as X grows very large or small. This problem is common to other mechanical methods of estimating volatility smiles such as splines.

One consequence is that the call delta is no longer a monotonically decreasing function of exercise price. The estimated implied volatility cannot be defined

EXHIBIT 7
CONTINUED

PANEL C. DOLLAR-YEN, DEALER I, MARCH 6, 1997

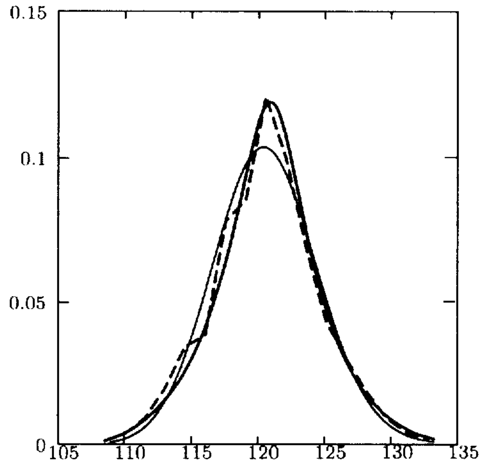

PANEL D. DOLLAR-YEN, DEALER II, MARCH 6, 1997

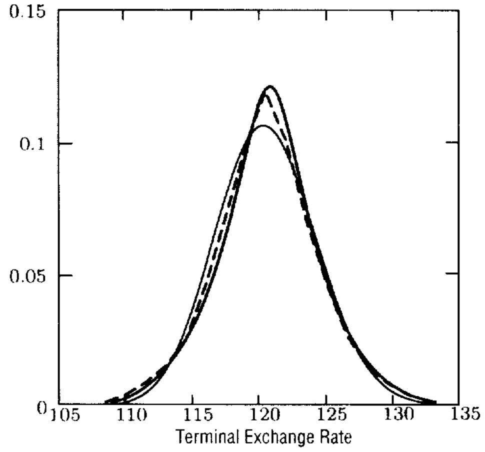

Heavy solid line: volatility function technique; dotted line: spline interpolation; thin solid line: lognormal approximation.

EXHIBIT 8
COMPARISON OF SHIMKO AND VOLATILITY FUNCTION TECHNIQUES

PANEL A. VOLATILITY SMILE: SHIMKO AND VOLATILITY FUNCTION INTERPOLATIONS

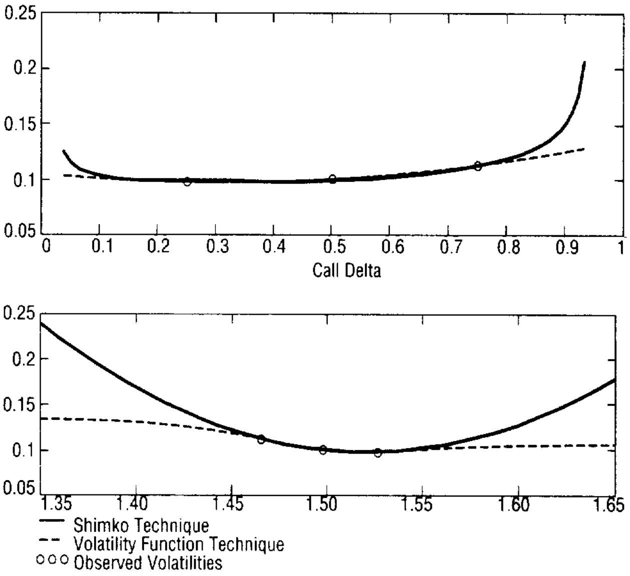

PANEL B. RISK-NEUTRAL DENSITY: SHIMKO AND VOLATILITY FUNCTION INTERPOLATIONS

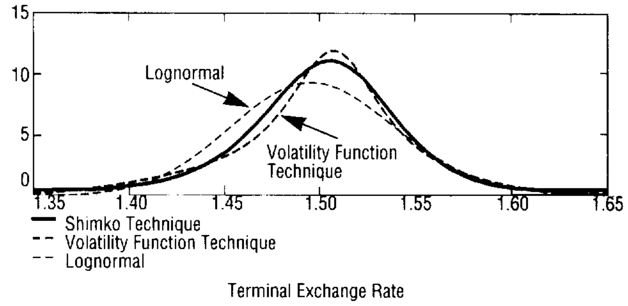

in terms of delta in the wings, since there is no unique exercise price associated with each low and high delta.

Another consequence is that the risk-neutral probability is not well-defined, since the second derivative of the call price with respect to the exercise price, with the implied volatilities supplied by the Shimko interpolation, does not integrate to zero ($e^{-r\tau}$) as the exercise price grows small (large). In the example in Exhibit 8, the graphs cover only the interval $\delta \in [0.03, 0.93]$ rather than $\delta \in [0, e^{-r^*\tau}]$. The implied volatility function avoids such violations, since it flattens out in the wings.[16]

To sum up, the volatility function method presented here generates highly accurate call prices for most exercise prices, although it may be off by small amounts for some very high or low exercise prices. Estimates of the risk-neutral density will thus be highly accurate except possibly for small errors in the extreme tails of the distribution.

As the over-counter currency options markets evolve, trading in farther out-of-the money calls and puts is likely to become more active. If dealers were to publicize prices of 10- or 15-delta risk reversals and strangles, one could calculate the implied volatilities of 10- and 90-delta calls, and use these to develop a version of the interpolation function $\hat{\sigma}_\delta(\delta)$ anchored at two additional points in the wings. Campa, Chang, and Reider [1997] give an example using weekly data over a one-year observation interval.

V. A SAMPLE APPLICATION

To see how this technique might be used by practitioners, let us examine an episode of dramatically changing sentiment: the dollar-yen exchange rate in mid-1995. From June 16, 1995, to August 16, 1995, spot dollar-yen rose from ¥84.50 to ¥98.15, and sentiment grew quite dollar bullish. The risk-neutral probability density functions of the exchange rate, displayed in Exhibit 9, indicate how much so. The x-axis measures the exchange rate as deviations from the contemporaneous forward dollar-yen rate, so the mean of each distribution is zero.

Uncertainty about the future exchange rate grew over the period, as can be seen from its greater dispersion and kurtosis on August 16. The salient feature of the densities is the change in skewness from markedly dollar bearish to dollar bullish. The risk-neutral likelihood of a sharp depreciation of the dollar was greater than that of an equally sharp appreciation on June 16, 1995, while the opposite was true on August 16, 1995.

Exhibit 10 displays the spot dollar-yen exchange rate and its coefficients of skewness and kurtosis, calculated from the one-month risk-neutral probability distribution, from April 1992 through early June 1996.

The forward premium, the risk-neutral expected value of the future exchange rate, offered little useful information on market sentiment about dollar-yen during the period that was not also embedded in the spot exchange rate; it certainly did not reveal the shift to dollar bullishness. The dollar's forward discount was unchanged at about –0.4%; current and anticipated central bank monetary operations were of subordinate importance in determining dollar-yen sentiment.

EXHIBIT 9
ESTIMATED RISK-NEUTRAL PROBABILITY DENSITY FUNCTION

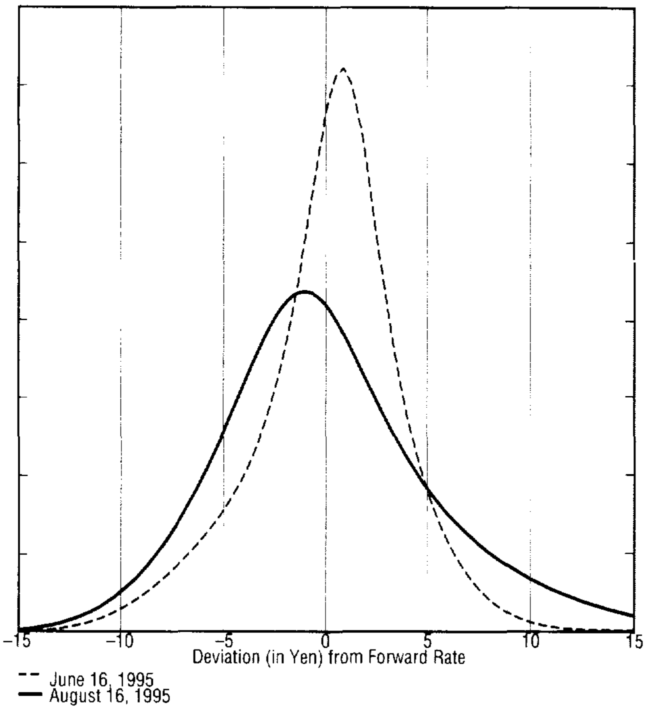

Dollar-yen, probability distribution of exchange rate one month hence.

It should be stressed that Exhibit 9 displays risk-neutral, not "true" subjective probability densities. Market participants will recall that in August 1995 substantial investor interest in out-of-the-money dollar calls (yen puts) materialized, motivated by expectations of substantial dollar appreciation over succeeding months. This left dealers short of volatility, quickly boosting option prices and driving risk reversal prices to unusually dollar-bullish levels.

It is worth keeping in mind that transient supply and demand imbalances, changes in market participants' exposures to the underlying currency, and changes in risk appetites may move option prices — and thus the risk-neutral distribution — without any change in expectations. However, as Exhibit 10 and the example of the dollar-yen in summer 1995 sug-

EXHIBIT 10

SKEWNESS AND KURTOSIS COEFFICIENTS OF THE DOLLAR-YEN EXCHANGE RATE — DAILY MARCH 31, 1992-JUNE 11, 1996

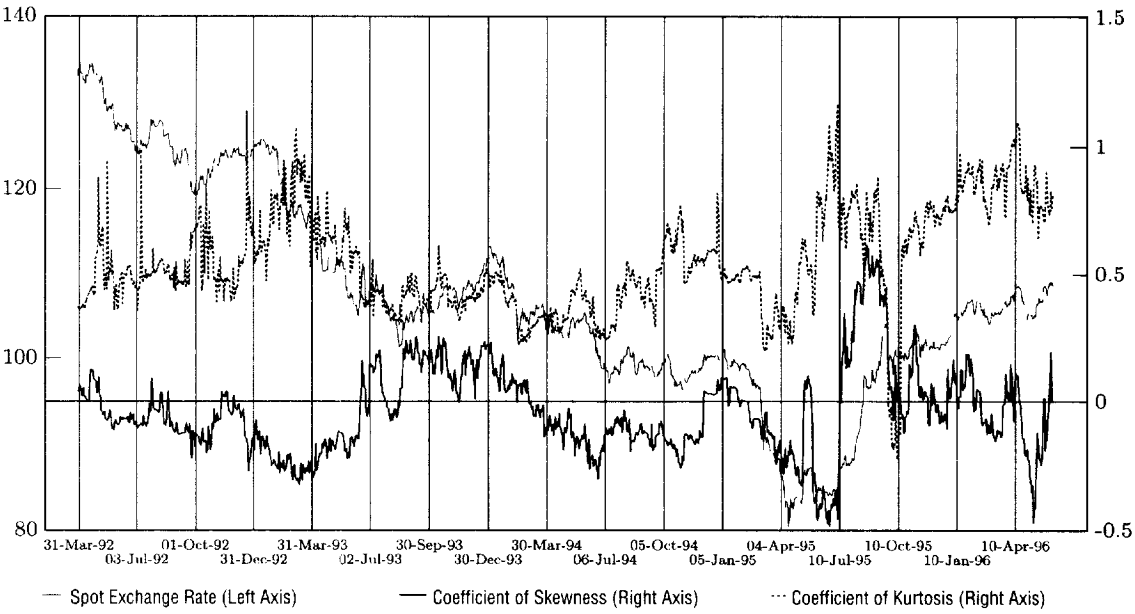

gest, changes in these factors in general coincide with changes in expectations, making it even more difficult to disentangle them.

Note that the risk-neutral skewness coefficient displayed in Exhibit 10 is biased toward the negative. This may be due to persistent dollar-bearish sentiment throughout the period, but since one observes negative skewness even during periods of dollar stability, it seems likely that the negative bias is due at least in part to the hedging requirements of Japanese exporters who persistently bid up the prices of out-of-the-money dollar puts.

Such biases may help account for the forward premium's failure as a predictor of changes in the spot rate. In fact, the risk-neutral moments derived from currency option prices — and the risk and expectational factors embedded in them — go a long way toward accounting for excess returns in exchange markets.[17]

VI. CONCLUSIONS

The method presented here makes it possible for market participants and observers to assess the risk-neutral probability distribution of exchange rates without expending substantial research and computer resources. To review its advantages:

- The data needed can be obtained by any market participant with foreign exchange dealing relationships or access to a Reuters terminal.
- The method is computationally simple and can be easily implemented on a PC. The estimate of the volatility smile is parameterized by the option prices, so it does not require the user to estimate a regression or spline function.
- The volatility function technique avoids the problems associated with the alternatives of constructing a spline or regressing the implied volatilities on delta or on exercise prices. It leads to smooth estimates of the risk-neutral density, avoiding the jagged appearance associated with these alternatives.
- It spite of its simplicity, it delivers a highly accurate representation of the risk-neutral distribution.
- It is consistent with the no-arbitrage conditions on risk-neutral distributions.

In all these respects, it compares favorably with alternative techniques.

Risk-neutral probability distributions calibrate for market participants what is meant by a large move in the exchange markets. Before seeing a risk-neutral distribution, many observers assign very high or low probabilities to, say, a 5% change in exchange rates. They also tend to overestimate the difference between the likelihood of an up or down move of equal size implied by risk reversal prices.

Looking at the risk-neutral distribution and its moments can help option dealers and investors use the information in option prices more effectively. Market participants can use the exchange rate's risk-neutral higher moments to quantify the riskiness of a currency position in a more differentiated fashion. In particular, they can distinguish the market's uncertainty about the exchange rate's future level — volatility risk — from the perceived risk of a particularly large exchange rate move, and they can also quantify the market view on exchange rate skewness. As we have seen, this information is a useful supplement to the generally rather uninformative forward exchange rate.

This technique can be of value to market observers and participants. For example, value at risk (VaR) techniques for the management of market risks often rely on the assumption that asset prices behave log-normally. One drawback of VaR is that it may inadequately account for the large but rare market moves that can potentially endanger a firm's very survival. Practitioners are aware of this, and devote much thought to guessing the degree of extra capital coverage required to compensate for the possibility of understating the likelihood of adverse "tail events." Risk-neutral probability distributions can serve as a useful complement to a variance-covariance matrix based on historical data.

APPENDIX A
RISK-NEUTRAL PROBABILITY DISTRIBUTION

We can think of the risk-neutral probability distribution as the probability density $\pi(x)$ that lets us represent observed call option prices $c(t, X, T)$ as the discounted expected values of their future payoffs:

$$c(t, X, T) = e^{-r\tau}E^*[\max(S_T - X)]$$

$$= e^{-r\tau}\int_X^\infty (S_T - X)\pi(S_T)dS_T \qquad \text{(A-1)}$$

where X is the exercise price, t and T are the current and option expiration dates, $\tau \equiv T - t$, r is the risk-free interest rate, E^* denotes an expectation taken under the risk-neutral probability distribution, S_t is the time t asset price, and $\int_a^b \pi(S_T)dS_T \equiv P^*(a \le S_T \le b)$, with P^* denoting a risk-neutral probability.

The risk-neutral distribution is *defined* as the probability measure — the set of numbers on [0, 1] integrating to unity — that makes Equation (A-1) "work"; that is, makes the expected return on the option equal the risk-free rate. The expected value of the future payoffs on any asset, calculated using these numbers as probabilities and discounted to the present at the risk-free rate, is equal to the current asset price.

To identify $\pi(x)$, we differentiate the call price with respect to the exercise price:

$$\frac{\partial c(t, X, T)}{\partial X} = -e^{-r\tau}[1 - \Pi(X)] \qquad \text{(A-2)}$$

where $\Pi(x) \equiv P^*(S_T \le x)$ is the risk-neutral cumulative distribution function, and

$$\frac{\partial^2 c(t, X, T)}{\partial X^2} = e^{-r\tau}\pi(X) \qquad \text{(A-3)}$$

A procedure to extract the probability distribution implied by a set of option prices treats $\pi(x)$ as an unknown function, which we try to estimate using Equation (A-3) and a set of observed option prices $c(t, X, T)$.

APPENDIX B
DERIVATION OF THE b_i

The at-the-money, or 50-delta volatility, sets $b_0 = 1$, since[18]

$$\text{atm}_t = \sigma_t^{(0.50)} = \sigma_\delta(0.50)$$

$$= b_0\text{atm}_t + b_1\text{rr}_t \cdot 0 + b_2\text{str}_t \cdot 0$$

Similarly, the risk reversal price sets $b_1 = -2$:

$$\text{rr}_t = \sigma_t^{(0.25)} - \sigma_t^{(0.75)}$$

$$= \sigma_\delta(0.25) - \sigma_\delta(0.75) = -\frac{b_1}{2}\text{rr}_t$$

Finally, the strangle price sets $b_2 = 16$, since

$$\text{str}_t = \frac{\sigma_t^{(0.75)} + \sigma_t^{(0.25)}}{2} - \text{atm}_t$$

$$= \frac{\sigma_\delta(0.25) + \sigma_\delta(0.75)}{2} - \text{atm}_t = 0.25^2 b_2 \text{str}_t$$

APPENDIX C

No-Arbitrage Conditions on the Volatility Function

The no-arbitrage condition on the derivative of market prices of European call options with respect to the exercise prices is derived by adding to Equation (A-2) the condition $0 \le \Pi(X) \le 1$. It can be stated as

$$-e^{-r\tau} \le \frac{\partial c(t, X, T)}{\partial X} \le 0 \qquad \text{(C-1)}$$

Condition (C-1) implies restrictions on the slope of the volatility function. We can express it equivalently in terms of implied volatilities using the volatility function:

$$-e^{-r\tau} \le \frac{\partial \hat{v}(X)}{\partial X} \equiv \frac{\partial v(S_t, \tau, X, \sigma_X(X), r, r^*)}{\partial X} \le 0 \quad \text{(C-2)}$$

Carrying out the differentiation in (C-2):

$$-e^{-r\tau} \le \frac{\partial \sigma_X(X)}{\partial(X)} \frac{\partial v(\cdot)}{\partial \sigma} - \frac{\partial v(\cdot)}{\partial X} \le 0$$

where $\partial v(\cdot)/\partial\sigma$ and $\partial v(\cdot)/\partial X$ refer to the partial derivatives of the Black-Scholes call pricing functions with respect to the implied volatility and exercise price.

The derivative of $\hat{v}(X)$ differs from the ordinary derivative of the Black-Scholes formula by the term $(\partial\sigma_X(X)/\partial X)$ $(\partial v(\cdot)/\partial\sigma)$, which reflects the distortion induced by the volatility smile. If there is no smile, that is, $\partial\sigma_X(X)/\partial X = 0$ everywhere, the no-arbitrage conditions are fulfilled.

The volatility function has modest curvature for intermediate exercise prices, and thus does not violate the no-arbitrage conditions in that range.[19] For high and low X, the volatility function's curvature tends to zero:

$$\lim_{X \to 0} \frac{\partial \sigma_X(X)}{\partial X} = \lim_{X \to \infty} \frac{\partial \sigma_X(X)}{\partial X} = 0$$

so the no-arbitrage conditions are preserved.

For large and small X, $\partial v(\cdot)/\partial\sigma$ tends toward zero, so the bounds on $\partial\sigma_X(X)/\partial X$ become quite loose. This explains the variation in simultaneous implied volatility quotes from dealers for options with deltas close to zero or $e^{-r^*\tau}$. A large change in the implied volatility produces only a small change in the option price in currency units.

It also shows that a good interpolating function for implied volatility need only be close to market-implied volatilities in the central range, say, deltas between 0.15 and 0.85. Outside that range, even fairly gross errors in the interpolation function will lead to small errors in call prices and thus in the estimated probability distribution (see Hodges [1996] for a more complete discussion, including a number of examples, of the no-arbitrage restrictions on volatility smiles).

ENDNOTES

The author acknowledges helpful discussions with Christophe Bristiel, Antoine Frachot, Dale Henderson, Karen Lewis, Rich Lyons, Jean McGillicuddy, Will Melick, Lars Olesen, Rob Reider, John Slee, Paul Söderlind, and Charlie Thomas; and participants at seminars at the Federal Reserve System, the International Monetary Fund, the Bank for International Settlements, Goldman, Sachs, and Co., Columbia University, and the 1996 Econometric Society European Meeting. James Robb and Tae Kim carried out the MATHLAB programming and calculations. The views expressed in this article are those of the author and do not necessarily reflect those of the Federal Reserve Bank of New York.

[1] See Appendix A for a formal presentation of this relationship, which was first noted by Breeden and Litzenberger [1978].

[2] See, for example, Shimko [1993]. Rubinstein [1994] and Derman and Kani [1994] estimate the entire stochastic process of the asset price using numerical methods.

[3] Examples of this approach include Bates [1996a, 1996b] and Malz [1996]. Melick and Thomas [1997] and Bahra [1996] present an estimation method in which the risk-neutral distribution is represented as a mixture of lognormals. While this approach is based on a distributional hypothesis, it is closer in spirit to interpolation, since the mixture distribution can represent a wide variety of distributions with a suitable choice of parameters.

[4] The application to options on currencies is often referred to as the Garman-Kohlhagen model, following its exposition in Garman and Kohlhagen [1983].

[5] To be precise, the expression in Equation (2) and here should involve the $75 - (1 - e^{-r^*\tau})$-delta, rather than

the 75-delta call volatility. For short-term options, the imprecision is unimportant unless the foreign interest rate is extremely high.

[6] Dealers often use shortcuts. They may use the strangle spread plus the at-the-money volatility in place of $\sigma_t^{(0.25)}$ and $\sigma_t^{(0.75)}$ to calculate the strangle price in currency units and use the at-the-money volatility plus or minus the risk reversal price in vols to calculate the risk reversal price in currency units. Typically, these shortcuts have an insignificant effect on the price in currency units.

[7] According to the most recent Bank for International Settlements survey, "Central Bank Survey" [1996], the over-the-counter market accounts for 97% of the notional outstanding stock of currency options.

[8] To be precise, the at-the-money volatility and risk reversal and strangle prices are available, but they can be easily transformed into the three volatility levels via Equation (8).

[9] One can use the same numerical techniques as in calculating implied volatilities, for example, MATLAB's solve function or Mathcad's root function.

[10] In the example, it is concave for deltas greater than about 0.85 and less than 0.15, about 4% or more above and below the current forward rate, as would be typical for flexible exchange rates with modest forward premiums and implied volatilities on the order of 10%.

[11] I have found for a variety of currencies using data since early 1992 that calculating $\hat{\Pi}(X)$ and $\hat{\pi}(X)$ over a range of exercise prices 25% above and below the current spot rate will almost always cover the entire one-month probability distribution. In general, plus or minus 10% covers virtually the entire distribution. Dividing that range into fifty steps will give adequate fineness. Some computational simplifications are described in Malz [1997].

[12] The risk-neutral first moment is $S_t e^{(r-r^*)(T-t)}$. The r-th central moment is defined by

$$\hat{\mu}_t^{(r)} = \int_0^\infty (X - F_{t,T})^r \hat{\pi}(X) dX$$

where $F_{t,T}$ is the forward rate. The risk-neutral annualized standard deviation of percent changes in the exchange rate is

$$\sqrt{\frac{\hat{\mu}_t^{(2)}}{\tau}}$$

and the risk-neutral coefficients of skewness and kurtosis are

$$\frac{\hat{\mu}_t^{(3)}}{[\hat{\mu}_t^{(2)}]^{3/2}}$$

and

$$\frac{\hat{\mu}_t^{(4)}}{[\hat{\mu}_t^{(2)}]^2} - 3$$

[13] Few dealers currently store historical data other than 25-, 50-, and 75-delta option volatilities, and those who do store only 10- or 15- and 85- or 90-delta volatilities, so most of these data are quite ephemeral, disappearing as they are updated.

[14] I use a fairly coarse partition mesh to avoid jagged probability densities. This problem is common to methods that do not impose a particular functional form on the volatility smile or on the risk-neutral distribution (see also Söderlind and Svensson [1997]).

[15] When applied to a larger data set, for example, settlement prices of exchange-traded options, Shimko's technique will not in general provide an exact fit to any actually observed option prices. The errors can be substantial, and if, as is often the case for exchange-traded options, most of the observations are near the money, the estimated implied volatilities of in- and out-of-the-money options may be highly sensitive to small changes in the data. Small errors, or even the tick size, may then have a significant impact on estimates.

[16] In fact, the integral is divergent. The reason is that the estimated implied volatilities rise so rapidly in the wings that the call deltas do not converge toward zero and one, but instead reverse direction, so that there are in fact two exercise prices for each low or high delta.

[17] Malz [1997] presents some empirical results on the power of risk-neutral distributions to explain currency excess returns and to predict future exchange rates.

[18] Recall that there is a very slight inaccuracy in setting $atm_t = \sigma_t^{(0.50)}$ and in using $\sigma_t^{(0.75)}$ rather than $\sigma_t^{[0.75-(1-e^{-r^*\tau})]}$.

[19] See Malz [1997] for details.

REFERENCES

Bahra, B. "Probability Distributions of Future Asset Prices Implied by Option Prices." *Bank of England Quarterly Bulletin*, 36, 3 (1996), pp. 299-311.

Bates, D.S. "Dollar Jump Fears, 1984-1992: Distributional Abnormalities Implicit in Currency Futures Options." *Journal of International Money and Finance*, 15, 1 (1996a), pp. 65-93.

——. "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options." *Journal of Financial Studies*, 9, 1 (1996b), pp. 69-107.

Breeden, D.T., and R.H. Litzenberger. "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of*

Business, 51, 4 (1978), pp. 621-651.

Campa, J.M., P.K. Chang, and R.L. Reider. "ERM Bandwiths for EMU and After: Evidence from Foreign Exchange Options." *Economic Policy*, 24 (1997), pp. 55-89.

"Central Bank Survey of Foreign Exchange and Derivatives Market Activity 1995." Basle: Bank for International Settlements, 1996.

Derman, E., and I. Kani. "Riding on a Smile." *Risk*, 7, 2 (1994), pp. 32-39.

Garman, M.B., and S.W. Kohlhagen. "Foreign Currency Option Values." *Journal of International Money and Finance*, 2, 3 (1983), pp. 231-237.

Heynen, R. "An Empirical Investigation of Observed Smile Patterns." *Review of Futures Markets*, 13, 2 (1994), pp. 317-354.

Hodges, H.M. "Arbitrage Bounds on the Implied Volatility Strike and Term Structures of European-Style Options." *Journal of Derivatives*, 3, 4 (1996), pp. 23-35.

Malz, A.M. "Option-Implied Probability Distributions and Currency Excess Returns." *Staff Report 32*: Federal Reserve Bank of New York, 1997.

——. "Using Option Prices to Estimate Realignment Probabilities in the European Monetary System: The Case of Sterling-Mark." *Journal of International Money and Finance*, 15, 5 (1996), pp. 717-748.

Melick, W.R., and C.P. Thomas. "Recovering an Asset's Implied PDF from Option Prices: An Application to Oil Prices During the Gulf Crisis." *Journal of Financial and Quantitative Analysis*, 32, 1 (1997), pp. 91-115.

Rubinstein, M. "Implied Binomial Trees." *Journal of Finance*, 49, 3 (1994), pp. 771-818.

Shimko, D. "Bounds of Probability." *Risk*, 6, 4 (1993), pp. 33-37.

Söderlind, P., and L.E. Svensson. "New Techniques to Extract Market Expectations from Financial Instruments." Discussion paper 1556, Center for Economic Policy Research, 1997.

Taylor, S.J., and X. Xu. "The Magnitude of Implied Volatility Smiles: Theory and Empirical Evidence for Exchange Rates." *Review of Futures Markets*, 13, 2 (1994), pp. 355-380.